arXiv:1609.04849v3 [stat.ML] 24 Feb 2017

Predicting Shot Making in Basketball Learnt from Adversarial Multiagent Trajectories

Mark Harmon [1] **Patrick Lucey** [2] **Diego Klabjan** [1]

Abstract

In this paper, we predict the likelihood of a player making a shot in basketball from multiagent trajectories. Previous approaches to similar problems center on hand-crafting features to capture domain specific knowledge. Although intuitive, recent work in deep learning has shown this approach is prone to missing important predictive features. To circumvent this issue, we present a convolutional neural network (CNN) approach where we initially represent the multiagent behavior as an image. To encode the adversarial nature of basketball, we use a multi-channel image which we then feed into a CNN. Additionally, to capture the temporal aspect of the trajectories we use "fading." We find that this approach is superior to a traditional FFN model. By using gradient ascent, we were able to discover what the CNN filters look for during training. Last, we find that a combined FFN+CNN is the best performing network with an error rate of 39%.

1. Introduction

Neural networks have been successfully implemented in a plethora of prediction tasks ranging from speech interpretation to facial recognition. Because of ground-breaking work in optimization techniques (such as batch normalization (Ioffe & Szegedy, 2015)) and model architecture (convolutional, deep belief, and LSTM networks), it is now tractable to use DNN methods to effectively learn a better feature representation compared to hand-crafted methods.

However, one area where such methods have not been utilized is the space of adversarial multiagent systems - specifically when the multiagent behavior comes in the form of trajectories. There are two reasons for this: i) procuring large volumes of data where deep methods are effective is difficult to obtain, and ii) forming an initial representation of the raw trajectories so that deep neural networks are effective is challenging. In this paper, we explore the effectiveness of deep neural networks on a large volume of basketball tracking data, which contains the x, y locations of multiple agents in an adversarial domain.

To throughly explore this problem, we focus on the following task: "given the trajectories of the players and ball in the previous five seconds, can we accurately predict the likelihood that a player with role X will make the shot?" We express this as a ten-class prediction problem by considering the underlying role of the player. For this paper, player role refers to the position of the player such as a point guard. Since we are utilizing an image-based technique, the CNN will not know which role shoots the ball, especially if there are multiple offensive players nearby. Thus, solving as a ten-class problem is ideal for our deep-learning approach.

Our work contains three main contributions. First, we create trajectories for the offense, ball, and defense as an eleven channel image. Each channel corresponds to the five offensive and defensive players, as well as the ball. In order to encode the direction of the trajectories, we fade the paths of the ball and players. In our case, an instance is a possession that results in an shot attempt. Second, we apply a combined convolutional neural network (CNN) and feed forward network (FFN) model on an adversarial multiagent trajectory based prediction problem. Third, we gain insight into the nature of shot positions and the importance of certain features in predicting whether a shot will result in a basket.

Our results show that it is possible to solve this problem with relative significance. The best performing model, the CNN+FFN model, obtains an error rate of 39%. In addition, it is able to accurately create heat maps by shot location for each player role. Features which are unsurprisingly important are the number of defenders around the shooter and location of the ball at the time of the shot. In addition, we found that our image-based model performs just as well as a feed-forward technique that required us to build nearly 200 features.

2. Related Work

With the rise of deep neural networks, sports prediction experts have new tools for analyzing players, match-ups, and team strategy in these adversarial multiagent systems.. Tra-

jectory data was not available at the time, so much previous work on basketball data using neural networks have used statistical features such as: the number of games won and the number of points scored. For example, Loeffelholz et al. (2009), use statistics from 620 NBA games and a neural network in order to predict the winner of a game. Another interested in predicting game outcomes is (McCabe & Trevathan, 2008). On the other hand, Nalisnick (2014) in his blog discusses predicting basketball shots based upon the type of shot (layups versus free throws and three-point shots) and where the ball was shot from. In other sports related papers, Huang & Chang (2010) use a neural network to predict winners of soccer games in the 2006 World Cup. Also, Wickramaratna et al. (2005) predict goal events in video footage of soccer games.

Although aforementioned basketball work did not have access to raw trajectory data, Lucey et al. (2014), use the same dataset provided by STATS for some of their work involving basketball. They explore how to get an open shot in basketball using trajectory data to find that the number of times defensive players swapped roles/positions was predictive of scoring. However, they explore open versus pressured shots (rather than shot making prediction), do not represent the data as an image, and do not implement neural networks for their findings. Other trajectory work includes using Conditional Random Fields to predict ball ownership from only player positions (Wei et al., 2015), as well as predicting the next action of the ball owner via pass, shot, or dribble (Yue et al., 2014). Miller et al. (2014) use non-negative matrix factorization to identify different types of shooters using trajectory data. Because of a lack of defensive statistics in the sport, A. Franks et. al. create counterpoints (defensive points) to better quantify defensive plays ((Franks et al., 2015b) (Franks et al., 2015a)). Perše et al. (2009) make use of trajectory data by segmenting a game of basketball into phases (offense, defense, and time-outs) to then analyze team behavior during these phases.

Wang & Zemel (2016) use trajectory data representations and recurrent neural networks (rather than CNN's) to predict plays. Because of the nature of our problem, predicting shot making at the time of the shot, there is not an obvious choice of labeling to use for a recurrent network. They also fade the trajectories as the players move through time. Similar to us, they create images of the trajectory data of the players on the court. Our images differ in that we train our network on the image of a five second play and entire possession, while their training set is based on individual frames represented as individual positions rather than full trajectories. They use the standard RGB channels, which we found is not as effective as mapping eleven channels to player roles and the ball for our proposed classification problem. Also, the images they create solely concentrate on offensive and do not include defensive positions.

In addition, work by Cervone et al. (2016) also explores shots made in basketball using the same STATS data. Our work focuses on a representation capturing the average adversarial multiagent behavior, compared to the approach of Cervone et.al., which focuses on individual players. In addition, while they concentrate on a Markov model that transitions between a coarse and fine data representation, our goal is to represent the fine-grained data as a whole for a deep learning model. Our work focuses on shot prediction for the average player in the NBA at the time of the shot while their work formulates a way to calculate estimated point value based upon specific player identity without reporting predictions for individual plays. Their estimated point value includes the probability of making a shot however this probability is based on an individual characteristics of a player taking the shot which is different from our starting point of considering generic players. We do not see an easy way to modify their models to provide answers based on our setting. In summary, there is no readily available numerical comparison of the two models that would provide answers for shot making predictions of generic players.

The final model that we implement, the combined network, utilizes both image and other statistical features. There is work that utilizes both image and text data with a combined model. Recently, Xu et al. (2015), Karpathy & Fei-Fei (2015), Socher et al. (2014), Vinyals et al. (2015), and Mao et al. (2015) all explore the idea of captioning images, which requires the use of generative models for text and a model for recognizing images. However, to the best of our knowledge, we have not seen visual data that incorporates a fading an entire trajectory for use in a CNN.

3. Data

The dataset was collected via SportsVU by STATS LLC. SportsVU is a tracking system that uses 6 cameras in order to track all player locations (including referees and the ball). The particular data used in this study was from the 2012-2013 NBA season and includes thirteen teams, which have approximately forty games each. Each game consists of at least four quarters, with a few containing overtime periods.

The SportVU system records the positions of the players, ball, and referees 25 times per second. At each recorded frame, the data contains the game time, the absolute time, player and team anonymized identification numbers, the location of all players given as (x,y) coordinates, the role of the player, and some event data (i.e. passes, shots made/missed, etc.). It also contains referee positions, which are unimportant for this study, and the three-dimensional ball location.

This dataset is very unique in that before SportVU, there was very little data available of player movements on the court and none known that provides frame-by-frame player locations. Since it is likely that most events in basketball can be determined by the movements of the players and the ball, having the trajectory data along with the event data should provide a powerful mixture of data types for prediction tasks.

There are a few ways to extract typical shot plays from the raw data. One is to choose a flat amount of time for a shooting possession. In our case we choose to include five seconds of a typical possession. To obtain clean plays, those that lasted less than 5 seconds due to possession changes and those in which multiple shots were taken were thrown out. After throwing out these cases, we were left with 75,000 five second plays.

The other way of obtaining play data would be to take the entire possession. Thus, rather than having plays be limited to five seconds, possessions can be much longer or shorter. Since the raw data does not contain labels for possession, we had to do this ourselves. To identify possession, we calculate the distance between the ball and each of the players. The player closest to the ball would be deemed the ball possessor. Since this approach may break during passes and other events during a game, we end the play of the possession when the ball is closer to the defensive team for 12 frames (roughly 0.5 seconds). The procedure yields 72,000 possession examples.

Since the classification problem is also dependent upon a player's role, a player's role must be chosen for each play. One way to do this would be to identify the role of a player at the beginning of the play and hold it constant. However, a player's role at the beginning of the play is usually not the same as a player's role at the time of the shot. Therefore, we decide to label a player's role via their position at end of the play in which the ball is shot.

4. Image-Based Representation

In terms of applying deep neural networks to multiagent trajectories, we first have to form an initial representation. A natural choice for representing these trajectories is in the form of an image. Given that the basketball court is 50x94 feet, we can form a 50x94 pixel image. In terms of the type of image we use, there are multiple choices: i) grayscale (where we identify that a player was a specific location by making that pixel location 1), ii) RGB (we can represent the home team trajectories in the red channel, the away team in the blue channel and the ball in the green channel, and the occurrence of a player/ball at that pixel location can be representing by a 1), and iii) 11-channel image (where each agent has their own separate channel). Examples of the grayscale and RGB approach are shown in Figure 1.

The 11-channel approach requires some type of alignment. In this paper, we apply the 'role-representation' which was first deployed by Lucey et al. (2013). The intuition behind this approach is that for each trajectory, the role of that player is known (i.e., point-guard, shooting guard, center, power-forward, small-forward). This is found by aligning to a pre-defined template which is learnt in the training phase.

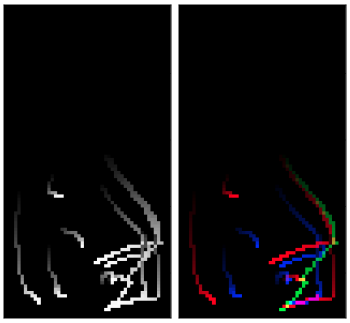

Figure 1. A Grayscale and RGB image of the same trajectory.

Figure 1 shows examples of the methods we use to represent our data for our CNN. The grayscale image, which appears on the left, can accurately depict the various trajectories in our system. However, because the image is grayscale, the CNN will treat each trajectory the same. Since we have an adversarial multiagent system in which defensive and offensive behavior in trajectory data can lead to different conclusions, grayscale is not the best option for representation. Therefore, to increase the distinction between our agents, we represent the trajectories with an RGB scale. We choose red to be offense, blue to be defense, and green to be the ball. This approach takes advantage of multiple channels to allow the model to better distinguish our adversarial agents. Although the ball may be part of the offensive agent structure, we decide to place the ball in a channel by itself since the ball is the most important agent. This approach, although better than the gray images, lacks in distinguishing player roles. Since we classify our made and missed shots along with the role of the player that shoots the ball, a CNN will have trouble distinguishing the different roles on the court. Therefore, for our final representation, we decide to separate all agents into their own channel so that each role is properly distinguished by their own channel.

The above ideas nearly creates ideal images; however, it does not include time during a play. Since each trajectory is of equal brightness from beginning to end, it may be difficult to identify where the ball was shot from and player locations at the end of the play. Therefore, we implement a fading variable at each time frame. We subtract a parame-

terized amount from each channel of the image to create a faded image as seen in Figures 1 and 2. Thus, it becomes trivial to distinguish the end of the possession from the beginning and leads to better model performance.

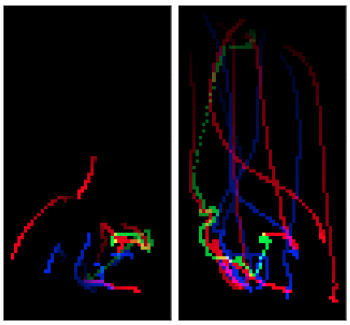

Figure 2. Figure on left depicts a five second play while the right is of an entire possession. As expected of five second plays, most of the trajectories remain near the basket located.

We create two options with our final faded image data: five seconds before the ball was shot and the total possession. Five seconds likely performs better since full possession data contains more temporal information than CNN's can handle. In longer possessions, trajectories tend to cross more, which confuses the CNN model by adversely affecting the fading since new trajectories will be higher in value.

5. Models

In order to fully utilize the power of this dataset, we implement a variety of networks for our prediction task. For our base model, we use logistic regression with 197 handcrafted features detailed later. To improve upon this basic model, we use a multi-layer FFN with the same features and utilize batch normalization during training. Because of the nature of these two models, we could only include the positions of the players at the time of the shot. Therefore, we craft images to include the position of the players throughout the possession. We then apply a CNN to these new image features. Finally, we create a combined model that adopts both images and the original FFN features for training.

5.1. Logistic Regression and Feed Forward Network

For the first models, features with the knowledge of the game in mind were crafted. The list of features includes:

- Player and ball positions at the time of the shot
- Game time and quarter time left
- Player speeds over either five seconds or an entire possession
- Distances and angles (with respect to the hoop) between players
- Whether the shooting player received a pass before the shot during the five second window
- Number of defenders in front of the shooter (30^0 of the shooter) and within six feet based upon the angles calculated between players
- Individual time of ball possession for each offensive player

Logistic regression and FFN both use the same calculated features. In addition, only the CNN does not incorporate the above features.

5.2. Convolutional Neural Network

For our model, we use a CNN that consists of three full convolutional layers, each with 32 3x3 filters and a max-pooling layer with a pool-size of 2x2 following each convolutional layer. After the last pooling layer there is a fully connected layer with 400 neurons, and finally an output layer consisting of a softmax function and ten neurons. In addition, we use the reLU function for our nonlinearity at each convolutional layer and the fully connected layer. We also implement AlexNet, Network-in-Network, and Residual Networks, but we did not garner significant improvement from either model.

5.3. CNN + FFN Network

The final network implemented is a combination of both the feed forward and convolutional networks. For this model, we use both the feed forward features and the fading trajectory images from the CNN. The idea behind the combined network is to have the model identify trajectory patterns in the images along with statistics that are known to be important for a typical basketball game.

The CNN and FFN parts of the combined network have the exact same architecture as the stand-alone versions of each model. The final layers just before the softmax layer of each stand-alone network are then fully-connected to a feed-forward layer that consists of 1,000 neurons. This layer is then fed into the final softmax layer to give predictions. After performing experiments and measuring log loss, we found that adding layers to this final network or adding additional neurons to this layer did not improve our final results.

6. Results

All of the models (FFN, CNN, and FFN + CNN) use the typical log loss function as the cost function with a soft-

max function at the output layer. The weights are initialized with a general rule of $\pm\sqrt{1/n}$ where n is the number of incoming units into the layer. For training, we implement the batch stochastic gradient method utilizing batch normalization. Batch normalization is used on the convolutional and feed forward layers of the model. In addition, we train the models on a NVIDIA Titan X using Theano.

For the CNN and CNN+FFN networks we utilize our eleven channel images with fading. To justify our representation, we predict our classification problem with each crafted image set with the previously mentioned CNN architecture. Figure 3 displays the log loss and error rate for each of our representations. Evaluating our representations by log loss and error rate show a dramatic difference between our images. The eleven channel method is understandably the best choice since having eleven channels gives the CNN much more role information.

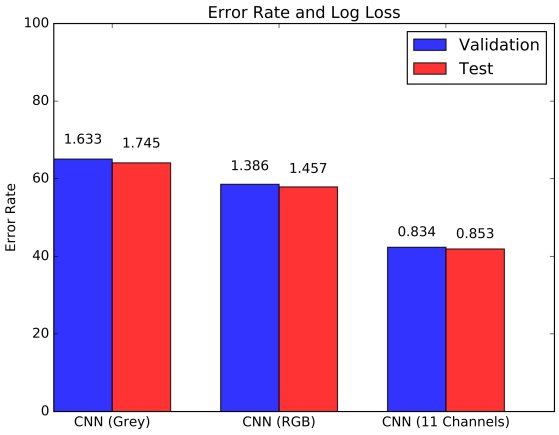

Figure 3. Results of Three Image Representations

Figure 4 gives a summary of the results of the three different models. It is intriguing that the convolutional neural networks captures the role behavior of our multiagent system without the addition of hand-crafted features that we use in the FFN. Note that this is only possible because we choose to represent the trajectories on eleven channels rather than three. Theoretically, the image data contains all pertinent information related to the play by only including positions, but it remains surprising that our CNN performs better than a traditional model with 200 features. Not only is making the images easier, but it saves time from having to come up with insightful features for input into a model such as a FFN.

The mode, Figure 5, expectedly surpasses the performance of either model separated, we can infer that the networks find different features during training. While the performance of the eleven channel images is an improvement (a

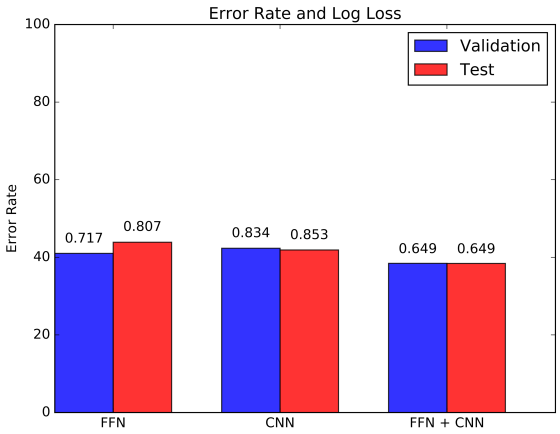

Figure 4. Final Results of All Three Models

10% gain), there is still much potential progress. Other methods include making each trajectory a different color in RGB space, varying the strength of the fading effect, and including extra channels of heat maps depicting the final ball position. Therefore, successfully representing trajectory data in sport that outperforms traditional metrics is a nontrivial problem and is not further explored here.

The remaining analyses are based on the combined model. In addition to assessing the accuracy of our model, we explore a basic heat map of basketball shots based upon the raw data. At the very least, we expect that our complete model should be able to recreate a heat map created via raw data. We make the heat map by taking a count of shots made against missed within a square foot of the basketball court. Since our classification model gives probabilities of making a shot (rather than a binary variable), we take the maximum probability to create a heat map equivalent to the raw data map.

In the raw data heat map, Figure 6, we note that the best probability of making a shot lies on top of the basket. As we get farther back, the probability decreases with two less probable zones (lighter zones separated by a thin green strip of higher probability): one right outside the paint and another just inside the three point line. The model results, Figure 5, expectedly prefers shots near the basket. The model also predicts a larger high value area surrounding the basket, which extends further into the paint of the court. However, the dead zones in the model heat map are much larger. In addition, the model over predicts near the basket while under predicting outside this area. Curiously, there are a few areas that have higher probability outside of the paint.

To further explore our results, we create heat maps solely based on the role of the player (to break down scoring

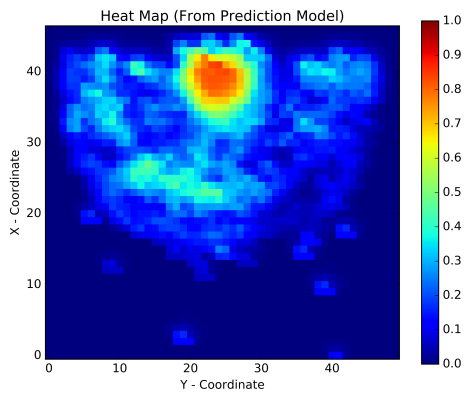

Figure 5. Heat Map from Model Data

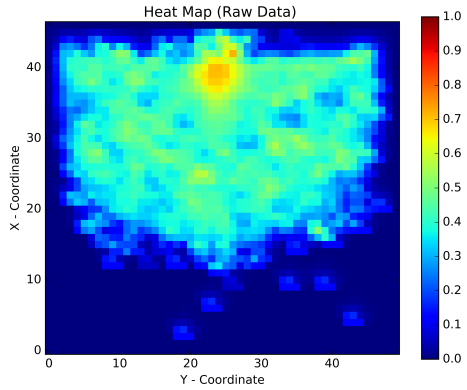

Figure 6. Heat Map from Raw Data

chances by agent). As before, each role represents an offensive player. In Figure 7 we present a few player roles and their representative heat maps. Role 3 is obviously the center position from their respective shot selection and Role 5 is the left guard. Note that the model predicts a much smaller area of midrange scoring probability than from the raw data for Role 3. The model heat map for Role 3 strictly covers the paint, while the raw data has significantly higher shot probabilities outside of the paint. Roles 1,4, and 5 show similar behavior in the model prediction. These maps are very heavy-handed with respect to under the basket shots with extremely small probabilities outside of the paint. The one exception is Role 2, which the model predicts has a much more likely chance of scoring outside of the paint. We observe from these break-down-by-role heat maps that Role 2 is the reason the heat map for all figures exhibits an arc of higher probability outside the hoop.

The probabilities that the combined model predicts for each shot may also provide useful insight into the game and our model's interpretation of high versus low value shots. We create these histograms by finding which examples the model gives the highest probability as a shot made or missed by player with role x. We then group these examples together, and the probability of making a shot is reported in the histogram. In the histogram Figure 8 we see that the majority of shots have a low probability. This agrees with common basketball knowledge because many of a guard's shots are beyond the paint. On the other hand, a center lives primarily under the basket and in the paint. Therefore, many of their shots are much more likely. Watching a game live, a center getting a clean pass right under the basket often results in a made shot. In addition, most roles tend to follow the probability pattern of Role 1 with the exception of Role 3 (the center), which has a wider distribution and higher average probability of making a shot. In addition, Role 5 tends to have a better ratio of high probability shots to low compared to Roles 1,2, and 4.

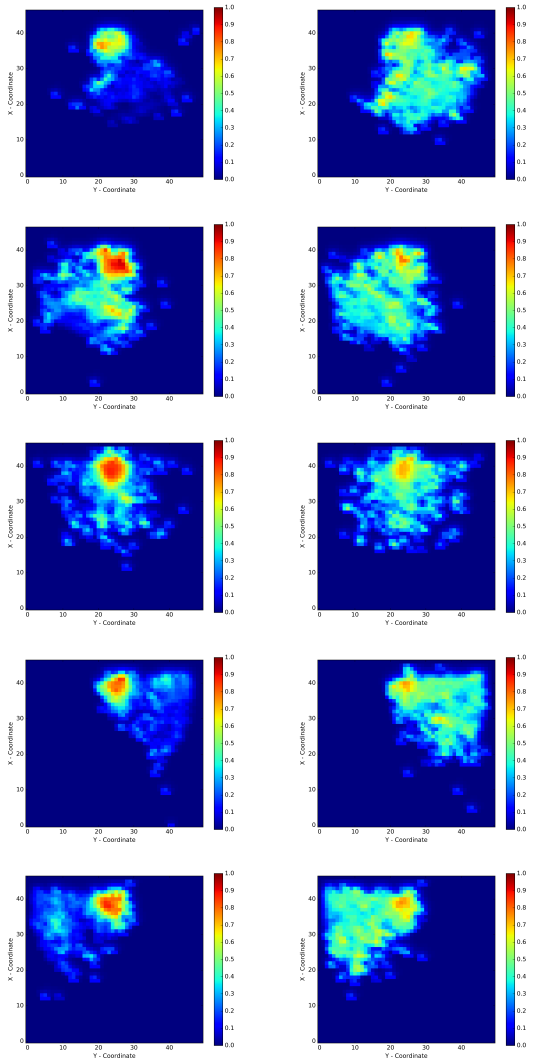

Figure 7. Top to bottom Roles 1-5. The left depicts the Model results and the right is the Raw Data results.

A brief glance at NBA statistics agrees with this interpretation as the players with the highest shooting percentage (barring free throws) tend to be centers.

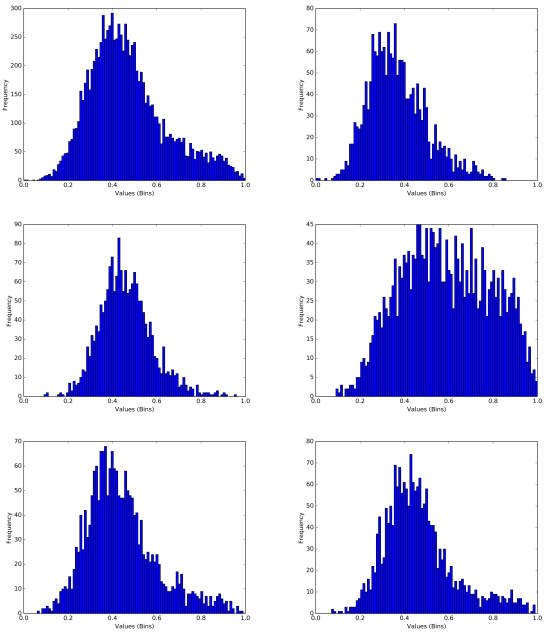

Figure 8. (1) All Roles. (2) Role 1. (3) Role 2. (4) Role 3. (5) Role 4. (6) Role 5.

When viewing the histograms more carefully, there are additional role dependent insights. For example, Role 1 has the lowest probabilities for shots made compared to all other roles. Although the general shape is the same as for Role 2, the predictions are shifted significantly to the left. We also note that overall, most of the shots fall into the 45% category of shot being made, which is aligned with general basketball knowledge. Since these histograms are predictions of the average basketball player in each role, the takeaway message is that unless you have an ace shooter (Lebron James or Steph Curry), a team should focus on ball movement to give the center an open shot rather than relying on outside shooters.

We also exhibit Figure 9 to provide additional visual context for our model probabilities. These figures depict the final positions of all players on the court at the time of the shot. We can assess how "open" the shot maker is at the time of the shot and the relative position of both the offensive and defensive players. The offensive players are blue, defensive players red, and the ball is green. Each offensive and defensive player has the letter "O" and "D" respectively followed by a number signifying the role of that player at that time. There are some times where the model makes some pretty questionable predictions. For example, the three-point shot that is exhibited in the top right of Figure 9 with a 0.518 probability is much too high. Unguarded, we do no expect three-point shots to be made more than 50% of the time. Thankfully, these examples are very rare in our model. For the most part, three-point shots are rated extremely low by the model garnering probabilities of less than 20%.

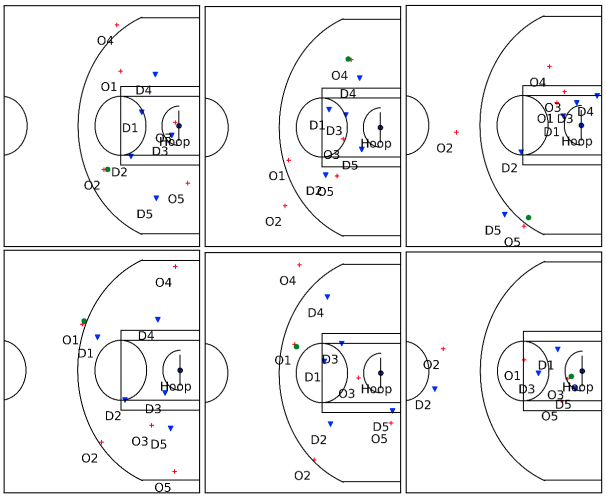

Figure 9. Role, Actual Shot, and Model Probability. (1) Role 2 Missed Shot (0.506). (2) Role 5 Made Shot (0.31). (3) Role 5 Missed Shot (0.518). (4) Role 1 Made Shot (0.392). (5) Role 1 Made Shot (0.245). (6) Role 3 Made Shot (0.857)

In addition to three-point shots having a generally low probability, shots that are well-covered by defenders have a much lower probability of success. This is an unsurprising well-known result, but it does add validity to our model. In addition, shots that are open and close to the basket are heavily favored in our model. For example, in the bottom right picture, Role 3 has a very good chance of making a wide open shot with the defense well out of position.

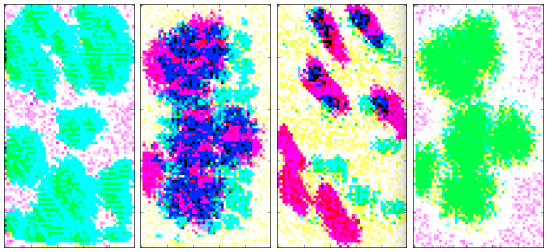

Figure 10. From left to right and top to bottom: (1) Filter 1 in Convolution Layer 1 (2) Filter 6 in Convolution Layer 3 (3) Filter 19 in Convolution Layer 19 (4) Filter 6 in Convolution Layer 1

Since our combined model predicts shots with more accuracy, we are curious to find the features that deep learning with a CNN observes with raw trajectory data. Our images contain both more temporal and spatial data since we fade images and do not include all multiagent interactions (since it is difficult to label all inter-player events). To investigate

our raw trajectory approach, we use a simple gradient ascent method to create an image that most strongly reacts to the filters in our network. For example, in a problem that contains cats and dogs, some filters may be looking for long whiskers to help identify a cat. In addition, these filters help determine whether the network has been well-trained. Networks having filters with little to no interpretation or appear random are usually poorly trained. Figure 10 shows a few representative images that we craft with the gradient ascent method.

In the first image of Figure 10 from the left, we see that the model is searching for defensive positioning (light blue) and the ball (green). There are clusters of these positions, which shows that the model is searching for the role that shoots the ball via ball positioning and the number of defensive players near the ball. On the other hand, the filter on the far right only depicts ball locations and does not include defensive positioning. The second figure from the left is actually very similar to the far right filter; however, it includes much more than simple ball information. This filter is specifically looking for ball location as well as offensive/defensive player locations near the ball. Thus, the CNN model is likely more effective in capturing defensive structure and paths for shot prediction, which is difficult to construct by hand.

7. Conclusions

Rather than other methods that incorporate transitions between a coarse and fine-grained approach or use small pieces of trajectory data, we utilize the full trajectory data of both offensive and defensive players. Since player alignments commonly permute, an image-based approach maintains the general spatial alignment of the players on the court. To integrate the time dependency of the trajectories, we introduce "fading" to our images to capture player paths. We found that using linear fading rather than a one-hot fade, works much better in our predictions. In addition, using a different color for each trajectory does not increase predictability in an RGB, three channel, image. Since traditional CNN's utilize three channels, we found that networks created to work with that kind of data, such as residual networks, AlexNet, and Network in Network underperformed. We therefore opt to build our own CNN to handle the trajectory images. Thus, by using our combined CNN and FFN, we can predict whether a shot is made with 61.5% accuracy with the CNN proving to be more accurate than a FFN with hand-crafted features.

We found that by using a CNN, we can further explore the data using gradient ascent to picture the various filters of our network. We found that as the network gets deeper, it tends to gather several features together. For example, in the first layer, the filters look for shot locations. As we delve deeper into the network, the filters also begin to look for defensive and offensive spatial positions to make a more accurate prediction. Also, the histograms agree with common knowledge that centers have the highest short percentage since their shots tend to be right beside the basket.

For further research, it would be very interesting to identify time dependency in basketball plays. In our image data, we subtract a flat amount at each equally spaced frame to cause the fading effect. However, this assumes that the data in time is linearly related. Since this is not necessarily true, designing a recurrent model to find this temporal dependency could be a very interesting problem. Instead of having a fading effect in the image data, we can design an LSTM that takes a moving window of player and ball trajectories.

One last aspect that was not taken into account during this study was the identities of teams and particular players. The focus of this research was to gather more insight on the average shooting plays of teams in the NBA. However, teams in the NBA have drastically different strategies. For example, the Golden State Warriors tend to rely on a three-point strategy while bigger teams, such the Thunder, build their offensive strategy around being inside the paint. Thus, new knowledge on basketball could be gathered if models were applied to different teams and possibly identify some overall team strategies. Such a more fine-grained analysis would require much more data.

References

Cervone, Daniel, DAmour, Alex, Bornn, Luke, and Goldsberry, Kirk. A multiresolution stochastic process model for predicting basketball possession outcomes. *Journal of the American Statistical Association*, 111(514):585–599, 2016.

Franks, Alexander, Miller, Andrew, Bornn, Luke, and Goldsberry, Kirk. Counterpoints: Advanced defensive metrics for nba basketball. In *MIT Sloan Sports Analytics Conference, Boston, MA*, 2015a.

Franks, Alexander, Miller, Andrew, Bornn, Luke, Goldsberry, Kirk, et al. Characterizing the spatial structure of defensive skill in professional basketball. *The Annals of Applied Statistics*, 9(1):94–121, 2015b.

Huang, Kou-Yuan and Chang, Wen-Lung. A neural network method for prediction of 2006 world cup football game. In *Neural Networks, The International Joint Conference on*, pp. 1–8. IEEE, 2010.

Ioffe, Sergey and Szegedy, Christian. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pp. 448–456, 2015.

Karpathy, Andrej and Fei-Fei, Li. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3128–3137, 2015.

Loeffelholz, Bernard, Bednar, Earl, Bauer, Kenneth W, et al. Predicting nba games using neural networks. *Journal of Quantitative Analysis in Sports*, 5(1):1–15, 2009.

Lucey, Patrick, Bialkowski, Alina, Carr, Peter, Morgan, Stuart, Matthews, Iain, and Sheikh, Yaser. Representing and discovering adversarial team behaviors using player roles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2706–2713, 2013.

Lucey, Patrick, Bialkowski, Alina, Carr, Peter, Yue, Yisong, and Matthews, Iain. How to get an open shot: analyzing team movement in basketball using tracking data. *MIT SSAC*, 2014.

Mao, Junhua, Xu, Wei, Yang, Yi, Wang, Jiang, Huang, Zhiheng, and Yuille, Alan. Deep captioning with multimodal recurrent neural networks (m-rnn). In *International Conference of Learning Recognition*, 2015.

McCabe, Alan and Trevathan, Jarrod. Artificial intelligence in sports prediction. In *Information Technology: New Generations, 2008. Fifth International Conference on*, pp. 1194–1197. IEEE, 2008.

Miller, Andrew, Bornn, Luke, Adams, Ryan P, and Goldsberry, Kirk. Factorized point process intensities: A spatial analysis of professional basketball. In *Internation Conference on Machine Learning*, pp. 235–243, 2014.

Nalisnick, E. Predicting basketball shot outcomes. `https://enalisnick.wordpress.com/2014/11/24/predicting-basketball-shot-outcomes/`, 2014.

Perše, Matej, Kristan, Matej, Kovačič, Stanislav, Vučkovič, Goran, and Perš, Janez. A trajectory-based analysis of coordinated team activity in a basketball game. *Computer Vision and Image Understanding*, 113(5):612–621, 2009.

Socher, Richard, Karpathy, Andrej, Le, Quoc V, Manning, Christopher D, and Ng, Andrew Y. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2:207–218, 2014.

Vinyals, Oriol, Toshev, Alexander, Bengio, Samy, and Erhan, Dumitru. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3156–3164, 2015.

Wang, K-c and Zemel, Richard. classifying nba offensive plays using neural networks. In *MIT Sloan Sports Analytics Conference, Boston, MA*, 2016.

Wei, Xinyu, Sha, Long, Lucey, Patrick, Carr, Peter, Sridharan, Sridha, and Matthews, Iain. Predicting ball ownership in basketball from a monocular view using only player trajectories. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 63–70, 2015.

Wickramaratna, Kasun, Chen, Min, Chen, Shu-Ching, and Shyu, Mei-Ling. Neural network based framework for goal event detection in soccer videos. In *Multimedia, IEEE International Symposium on*, pp. 8–pp. IEEE, 2005.

Xu, Kelvin, Ba, Jimmy, Kiros, Ryan, Cho, Kyunghyun, Courville, Aaron C, Salakhutdinov, Ruslan, Zemel, Richard S, and Bengio, Yoshua. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, volume 14, pp. 77–81, 2015.

Yue, Yisong, Lucey, Patrick, Carr, Peter, Bialkowski, Alina, and Matthews, Iain. Learning fine-grained spatial models for dynamic sports play prediction. In *Data Mining, 2014 IEEE International Conference on*, pp. 670–679. IEEE, 2014.